JOSHUA ERMERT

San Diego, CA 92120 ▪ 760-688-9785 ▪ jseermert@gmail.com
linkedin.com/in/josh-ermert-79496b176 ▪ github.com/JSErmert

PROFILE

AI engineer focused on responsible, governed AI: human-in-the-loop systems, verified source-grounding, and anti-confabulation discipline. Builds generative-AI applications and, through natural-language extraction, queryable knowledge libraries that turn unstructured archives into searchable, source-verified assets, with machine-learning grounding in Python and R. Strong with GitHub-based workflows, technical documentation, and independent remote work.

EDUCATION

SAN DIEGO STATE UNIVERSITY: B.S. Management Information Systems · Double Minor in Computer Science & Interdisciplinary Studies · Weber Honors College · National Society of Leadership and Success · GPA 3.52 · Cum Laude *May 2026*

UNIVERSITY OF SAN DIEGO: GenCyber Cybersecurity Summer Program · NSA / NSF-funded academic cybersecurity camp · defensive security primitives · threat modeling · secure system administration *May 2019*

Relevant coursework · Reinforcement Learning (CS 654, A, graduate-level) · Machine Learning (CS 549, A) · Linear Algebra (MATH 254, A) · Programming Languages (CS 320, A) · Data Structures (CS 310, C++) · Computer Architecture (CS 370, MIPS Assembly) · Calculus I & II · Physics I & II (with labs) · Statistics (STAT 119, A) · Digital Entrepreneurship & Innovation (MIS 455, A)

PORTFOLIO PROJECTS

VisionAir: AI Project-Genesis & Context Engine for Agentic AI *github.com/JSErmert/visionair · visionair-sable.vercel.app*

> *Live web app that turns an ambiguous idea into a structured, ready-to-build context pack for AI coding agents. A coverage-driven adaptive interview asks only about what is missing, then synthesizes and assembles a downloadable, governed project pack.*

Solo · Next.js 16 · TypeScript · React 19 · Anthropic Claude (Sonnet 4.6 + Opus 4.7) · Postgres · Vercel · GitHub Actions CI/CD

- Built a generative project-genesis pipeline (adaptive interview plus synthesis) that turns an unstructured idea into a structured, documented, ready-to-build asset, with versioned workspaces and incremental refinement.
- Designed the relational data layer on Postgres (schema, migrations, per-owner queries) so every generated artifact persists and is retrievable.
- Hardened the AI surface with prompt-injection defenses, a deterministic fallback on rate limits, and CI/CD security scanning.

Diabetes Risk Prediction by Patient Health Indicators *github.com/JSErmert/AEI-POC*

> *End-to-end statistical analysis pipeline on the CDC BRFSS 2015 survey dataset (253,680 records, 22 features), fully reproducible from data ingestion through model comparison. Produced as the MIS 401 final project.*

Solo · R · Python · pandas · numpy · matplotlib · python-pptx

- Trained and compared multinomial logistic regression, k-tuned KNN, and a pruned decision tree on a fixed 70/30 split (seed 401) for full reproducibility, then benchmarked their predictive performance.
- Built a Python pipeline that produced the reporting deck programmatically from the same source data the modeling code used, preserving evidence lineage from raw data to the final artifact.

HydrOS: Reinforcement-Learning Simulation for a Microplastics-Extraction System *github.com/JSErmert/hydrion*

> *A multi-layer reinforcement-learning simulation environment for a microplastics-extraction system, built to model and train control policies for constrained system behavior under stress before any hardware build. Primarily back-end systems engineering and reinforcement learning.*

Solo · Python · PPO / Stable Baselines3 · NumPy · multi-physics simulation

- Architected a multi-physics simulation environment and trained PPO agents under immutable, milestone-based execution contracts.
- Designed validation-driven testing across reward shape, conservation invariants, and per-component physics, trusting a design only after it survives adversarial validation.

PROFESSIONAL EXPERIENCE

INDEPENDENT AI CONSULTANT *2025 – Present*

> *Governed, source-grounded AI knowledge systems for private clients (paid) · system architecture is my IP*

- For an independent physical-therapy professional, designed and built a governed-AI validated research library that turns unstructured professional expertise into searchable, source-verified, traceable records, with a human reviewer in control of every decision.
- Built a source-grounding layer that verifies every claim against the published literature and refuses to fabricate; across a 30-day production run it grounded 100+ deliverables with 28 source-verified citations and zero fabricated references.

- For a small-business client, built a governed operations system: a layered private corpus surfaced through a routed assistant that recommends actions while a human commits every change, never deploying autonomously.
- Client domain content is each client's IP and is not shown; the system architecture is my IP.

CELLTRION USA: Commercial Operations, Data Analytics & AI Solutions Intern *6/2025 – 8/2025*

- Developed an Azure OpenAI integration architecture and implementation plan for AI-assisted operational analytics, covering secret handling, prompt patterns, and human-in-the-loop review gating.
- Stood up anomaly detection over 10,000+ commercial-operations records with structured review workflows, including ML-classification feature design for an internal analytics initiative.
- Built two production Power BI dashboards used weekly by the full Sales and Commercial Operations teams, with territory alignment spanning all 50 states and U.S. territories.
- Collaborated cross-functionally with Sales, Operations, and IT stakeholders to validate output accuracy and resolve data discrepancies between source systems.

IPERDESIGN, Rome, Italy: UX / UI Design Intern (250 hours) *6/2023 – 7/2023*

- Designed and built 3 client websites in WordPress and Figma over a 250-hour international internship, working closely with stakeholders across language audiences.
- Reviewed and updated English-language content for Iperdesign's own site and a client presentation, applying editorial QA to preserve message clarity and descriptive accuracy.

CITY OF SAN DIEGO: Cyber Security Intern *7/2022 – 9/2022*

- Completed UCSD Linux lab exercises and cybersecurity fundamentals training, building hands-on grounding in command-line system administration and networking fundamentals.
- Coordinated learning operations across the intern cohort, organizing mentorship sessions and weekly technical activities, including explaining technical concepts clearly to non-technical learners.

TECHNICAL SKILLS

AI & Machine Learning • generative AI (Anthropic Claude, Azure OpenAI) · prompt engineering · context engineering · natural language processing (LLM-based extraction) · reinforcement learning (PPO / Stable Baselines3) · classification (logistic regression, KNN, decision trees) · Python · R
Responsible & Governed AI • human-in-the-loop review gating · verified source-grounding (anti-fabrication, exact-figure citation) · AI output evaluation (accuracy, fabrication, integrity) · validation discipline · prompt-injection defenses · audience-aware output governance · evidence lineage
Data & Knowledge Structuring • searchable / queryable knowledge libraries · unstructured-to-structured corpus extraction · metadata management, schemas & taxonomy · controlled vocabulary · index & bidirectional traceability · archival preparation · SQL · Power BI · Excel · pandas
Engineering & Tooling • Python · TypeScript / JavaScript · React · Next.js · Node.js · Git / GitHub · GitHub Actions CI/CD · technical writing & documentation
Foundations & Security • GenCyber (NSA / NSF) · secrets management · dependency scanning (Trivy, Dependabot) · TLS / security headers · networking fundamentals (TCP/IP, DNS)